

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Methods and Systems for Protecting Information in  
Paging Operating Systems**

Inventor(s):  
Scott A. Field

ATTORNEY'S DOCKET NO. MS1-407US

1 **PRIORITY APPLICATION**

2 This application stems from and claims priority to U.S. Provisional Patent  
3 Application No. 60/143,438, filed on July 13, 1999, the disclosure of which is  
4 hereby incorporated by reference.

5  
6 **TECHNICAL FIELD**

7 This invention relates to paging operating systems and, more particularly, to  
8 methods and systems of protecting information within paging operating systems.

9  
10 **BACKGROUND**

11 Computer systems, in general, include a main memory (also known as the  
12 computer's "physical memory") for storing data and instructions of currently  
13 executing programs ("process threads"). Typically, the main memory is organized  
14 as a plurality of sequentially numbered storage units, each containing a fixed size  
15 quantity (e.g. an 8-bit byte in byte oriented computers). The numbering of the  
16 storage units (typically in binary or hexadecimal values starting from zero up to  
17 the total number of storage units minus one) serve as addresses by which a  
18 particular storage unit can be referenced for reading or writing the data contained  
19 therein. The set of numbers by which the storage units are addressed is known as  
20 the "physical address space" of the main memory. Main memory typically is  
21 realized using semiconductor memory which provides fast, random-access to the  
22 various storage units, but requires constant application of electrical energy for  
23 operation (i.e. the memory is volatile).

24 Computer systems also typically provide one or more secondary storage or  
25 memory devices which are generally slower than the main memory, but have a

much greater storage capacity than the main memory. The secondary storage devices typically store data on a magnetic or optical media that is non-volatile, such as a hard disk. Secondary storage devices generally store data in the form of files or sequential data streams.

Due to the greater speed at which data can be accessed in main memory, data that is currently in use by process threads running on the computer system is desirably stored in the main memory. Due to the smaller storage capacity of the main memory, however, main memory may be unable to store all the information needed by process threads. Accordingly, data that is no longer currently in use is desirably removed from the main memory, or moved from the main memory to the secondary storage devices.

Techniques to efficiently manage the use of the main memory ("memory management techniques") by process threads are conventionally known. One standard technique, commonly known as "virtual memory," is implemented by many operating systems, usually in cooperation with a computer system's processor. Virtual memory techniques create a separate address space, referred to as the "virtual address space" or "process address space" by which process threads access data in memory. The operating system and processor translates or maps a subset of the virtual addresses in the virtual address space to actual physical addresses in the main memory's physical address space. When a process thread reads or writes data to a virtual address in its virtual address space, the operating system and/or processor translates the virtual address to a corresponding physical address of a storage unit in the main memory where the data is to be read or written. In Microsoft Corporation's Windows NT operating system, for example, a

1 component called the virtual memory manager implements a separate virtual  
2 address space for each process in cooperation with the computer's processor.

3 Since the virtual address space is typically much larger than the physical  
4 address space of the main memory, only a subset of the virtual address space can  
5 be resident in main memory at one time. Data not resident in main memory is  
6 temporarily stored in a "backing store" or "paging" file on the computer's hard  
7 disk. When the main memory becomes over committed (i.e. its storage capacity is  
8 exceeded), the operating system begins swapping some of the contents of the main  
9 memory to the "backing store" file. When the data is again required by a process  
10 thread, the operating system transfers the data back into the main memory from  
11 the backing store file. By swapping data that is no longer needed to the hard disk,  
12 virtual memory allows programmers to create and run programs that require more  
13 storage capacity than is available in the main memory alone.

14 Moving data between the main memory and the hard disk is most  
15 efficiently performed in larger size blocks (as compared to bytes or words).  
16 Accordingly, virtual memory techniques generally perform swapping in large size  
17 blocks. Microsoft Corporation's Windows NT operating system, for example,  
18 divides the virtual address space of each process thread into equal size blocks  
19 referred to as "pages." The main memory also is divided into similar size blocks  
20 called "page frames," which contain the pages mapped into the main memory.  
21 The page size in the Windows NT operating system can vary depending on the  
22 requirements of the particular computer on which it is run.

23 In the Windows NT operating system, each process has a set of pages from  
24 its virtual address space that are present in physical memory at any given time.  
25 Pages that are currently in the main memory and immediately available are termed

1 "valid pages." Pages that are stored on disk (or in memory but not immediately  
2 available) are called "invalid pages." When an executing thread accesses a virtual  
3 address in a page marked "invalid", the processor issues a system trap called a  
4 "page fault." The operating system then locates the required page on the hard disk  
5 and loads it into a free page frame in the main memory. When the number of  
6 available page frames runs low, the virtual memory system selects page frames to  
7 free and copies their contents to the hard disk. This activity, known as "paging," is  
8 imperceptible to the programmer.

9 One of the problems that continues to confront so-called paging operating  
10 systems, such as the one described above, concerns the treatment of sensitive  
11 information (e.g. passwords to access network resources, credit card information  
12 used during an Internet shopping session, and the like). For example, when an  
13 individual, using a password, logs onto an operating system such as Windows NT,  
14 the individual's password can typically be kept in memory for various reasons.  
15 For example, if the user locks a work station and wants to later unlock it, the  
16 operating system needs to validate against something. Thus, the operating system  
17 goes out to main memory and compares what is typed in by a user with what is  
18 sitting in the memory. Between these two points in time, however, the password  
19 may have entered the paging file because the operating system may have decided  
20 that the logon process was idle. Having the password in the paging file can leave  
21 it open to attack, e.g. if the machine on which the paging file is located were to be  
22 physically stolen. Thus, because of the nature of paging operating systems,  
23 sensitive information can sometimes be undesirably placed in a paging file in  
24 secondary memory. In security-sensitive installations, preventing the sensitive  
25 information from reaching the paging file may be advantageous.

1        There have been attempts in the past to address the situation of sensitive  
2 information making it into the paging file. These attempts have been successful in  
3 some respects, but still fall short of the mark insofar as providing a system that is  
4 desirably secure and economical to use.

5        One past approach has been to designate certain pages of the main memory  
6 as "page locked," and to place sensitive information only in page-locked pages.  
7 The "page locked" designation is a flag that tells the memory manager that the  
8 designated page is never to be moved to the paging file. While this ensures that  
9 the sensitive information does not make its way to the paging file, it consumes  
10 valuable main memory. Because there is a finite amount of main memory  
11 available, this approach is not optimal.

12        Another approach has been to configure the operating system to zero  
13 portions of the page file which are no longer associated with allocated memory  
14 when the operating system is shut down. This approach is problematic in the  
15 event that any of the following two events occurs: (1) power loss –induced or  
16 accidental, and (2) pages of memory are still allocated and active that contain  
17 sensitive information. In the former case, once power loss has occurred, an  
18 attacker can analyze the page file and "undo" any obfuscation as necessary. The  
19 operating system never had a change to zero the page file which would normally  
20 occur during clean shutdown of the operating system. In the latter case, if the  
21 pages of memory are still allocated and active, the operating system will be unable  
22 to zero the sensitive information that is contained in such pages.

23        Yet another approach has been to encrypt the sensitive information with a  
24 key that is hard-coded somewhere in the operating system. When the sensitive  
25 information is then sent to the hard disk, it will be encrypted and theoretically safe.

1 This approach is not optimal because it is still subject to attack. Specifically, an  
2 attacker who accesses the hard disk need only look for data that appears to have  
3 been obfuscated with a key and then set about to break the key. Obfuscated or  
4 encrypted information can be recognized using a variety of approaches, for  
5 example, measuring the entropy of blocks of data. Once the key is broken, all of  
6 the encrypted information can be accessed. And, because the key is hard-coded, it  
7 never changes. Thus, once it is discovered by an attacker, the attacker can have  
8 access to all information that has been or will be encrypted using the hard-coded  
9 key.

10 This invention arose out of concerns associated with providing improved  
11 methods and systems for protecting information that is used in paging operating  
12 systems.

### 13 14 SUMMARY

15 The inventive methods and systems provide an approach to protecting  
16 unencrypted sensitive information from being paged out to secondary memory,  
17 such as a hard disk, during paging operations. In the described embodiment, a key  
18 is provided and is maintained in the main memory of a virtual memory system.  
19 Measures are taken to protect the key such as page-locking the key in the main  
20 memory to ensure that it never gets paged out to the secondary memory. The  
21 illustrated key is a desirably large key that is randomly generated by the operating  
22 system. When sensitive information is to be placed in the main memory, it is  
23 encrypted with the page-locked key. The encrypted sensitive information can then  
24 be paged out to secondary memory without concern about its security. When the  
25 encrypted sensitive information is needed by a process or application, it is

retrieved from secondary memory and decrypted using the page-locked key. For further protection, the sensitive information can be decrypted into a page-locked page of main memory. More than one key can be used to encrypt and/or decrypt the sensitive information.

In one aspect, the encryption/decryption process can be initiated by one or more applications. The applications initiate the encryption process by calling a software component that handles the encryption/decryption. In the described embodiment, the software component comprises the operating system kernel. The software component retrieves the page-locked key and performs the encryption on the sensitive information.

In another aspect, the memory manager is closely integrated with the encryption/decryption process. Specifically, each page in main memory has an attribute that can indicate that it is a secure page. When the memory manager handles secure pages, it is programmed to ensure that the page is encrypted before paging it out to the paging file. All encryption/decryption thus takes place through the memory manager.

## **BRIEF DESCRIPTION OF THE DRAWINGS**

Fig. 1 is an exemplary computer system that can be used to implement the described embodiments.

Fig. 2 is a block diagram of an operating system architecture within a computer system such as the one shown in Fig. 1.

Fig. 3 is a diagrammatic representation of an exemplary initiation/protection phase in accordance with a described embodiment.



1 Fig. 4 is a flow diagram that describes steps in a method in accordance with  
2 a described embodiment.

3 Fig. 5 is a block diagram of internal functions of the operating system  
4 kernel and virtual memory manager for protecting and unprotecting information.

5 Fig. 6 is a flow diagram that describes steps in a method in accordance with  
6 a described embodiment.

7 Fig. 7 is a flow diagram that describes steps in a method in accordance with  
8 a described embodiment.

9 Fig. 8 is a diagrammatic representation of an exemplary access phase in  
10 accordance with a described embodiment.

11 Fig. 9 is a flow diagram that describes steps in a method in accordance with  
12 a described embodiment.

## 13 14 **DETAILED DESCRIPTION**

### 15 **Exemplary Computer System**

16 Fig. 1 shows a general example of a desktop computer 130 that can be used  
17 in accordance with the invention. Computer 130 includes one or more processors  
18 or processing units 132, a system memory 134, and a bus 136 that couples various  
19 system components including the system memory 134 to processors 132. The bus  
20 136 represents one or more of any of several types of bus structures, including a  
21 memory bus or memory controller, a peripheral bus, an accelerated graphics port,  
22 and a processor or local bus using any of a variety of bus architectures. The  
23 system memory 134 includes read only memory (ROM) 138 and random access  
24 memory (RAM) 140. A basic input/output system (BIOS) 142, containing the  
25



1 devices are connected to the processing unit 132 through an interface 170 that is  
2 coupled to the bus 136. A monitor 172 or other type of display device is also  
3 connected to the bus 136 via an interface, such as a video adapter 174. In addition  
4 to the monitor, personal computers typically include other peripheral output  
5 devices (not shown) such as speakers and printers.

6 Computer 130 commonly operates in a networked environment using  
7 logical connections to one or more remote computers, such as a remote computer  
8 176. The remote computer 176 may be another personal computer, a server, a  
9 router, a network PC, a peer device or other common network node, and typically  
10 includes many or all of the elements described above relative to computer 130,  
11 although only a memory storage device 178 has been illustrated in Fig. 1. The  
12 logical connections depicted in Fig. 1 include a local area network (LAN) 180 and  
13 a wide area network (WAN) 182. Such networking environments are  
14 commonplace in offices, enterprise-wide computer networks, intranets, and the  
15 Internet.

16 When used in a LAN networking environment, computer 130 is connected  
17 to the local network 180 through a network interface or adapter 184. When used  
18 in a WAN networking environment, computer 130 typically includes a modem 186  
19 or other means for establishing communications over the wide area network 182,  
20 such as the Internet. The modem 186, which may be internal or external, is  
21 connected to the bus 136 via a serial port interface 156. In a networked  
22 environment, program modules depicted relative to the personal computer 130, or  
23 portions thereof, may be stored in the remote memory storage device. It will be  
24 appreciated that the network connections shown are exemplary and other means of  
25 establishing a communications link between the computers may be used.

1 Generally, the data processors of computer 130 are programmed by means  
2 of instructions stored at different times in the various computer-readable storage  
3 media of the computer. Programs and operating systems are typically distributed,  
4 for example, on floppy disks or CD-ROMs. From there, they are installed or  
5 loaded into the secondary memory of a computer. At execution, they are loaded at  
6 least partially into the computer's primary electronic memory. The invention  
7 described herein includes these and other various types of computer-readable  
8 storage media when such media contain instructions or programs for implementing  
9 the steps described below in conjunction with a microprocessor or other data  
10 processor. The invention also includes the computer itself when programmed  
11 according to the methods and techniques described below.

12 For purposes of illustration, programs and other executable program  
13 components such as the operating system are illustrated herein as discrete blocks,  
14 although it is recognized that such programs and components reside at various  
15 times in different storage components of the computer, and are executed by the  
16 data processor(s) of the computer.

### 17 18 **Exemplary Operating System**

19 The described embodiment is illustrated in the context of Microsoft  
20 Corporation's Windows NT operating system. For a complete discussion of the  
21 Windows NT operating system see "Inside Windows NT", by Helen Custer,  
22 Microsoft Press, 1993. A brief overview of part of the general operation of the  
23 Windows NT operating system and associated terminology is provided below  
24 before discussing the details of the described embodiment.  
25



1 respectively) run, with a limited set of interfaces available and with limited access  
2 to system data.

3 Each subsystem implements a single set of services, for example, memory  
4 services, process creation services, or processor scheduling services. The  
5 illustrated Win32 subsystem 68, for example, makes a 32-bit application  
6 programming interface (API) available to application programs. Each subsystem  
7 runs in user mode, executing a processing loop that checks whether a client 72 has  
8 requested one of its services. The client 72 may be another operating system  
9 component or an application program.

10 The executive 60 is a series of components, each of which implements two  
11 sets of functions: system services 74, which can be called from client programs as  
12 well as other executive components, and internal routines, which are available  
13 only to components within the executive. System services include (a) the object  
14 manager 60a, which is responsible for creating, managing and deleting objects  
15 (objects are abstract data structures used to represent operating system resources);  
16 (b) the process manager 60b, which is responsible for creating/terminating  
17 processes and threads, and for suspending/resuming execution of threads; and (c)  
18 the I/O manager 60f, which is responsible for implementing device-independent  
19 I/O facilities as well as device-dependent I/O facilities.

20 The client 72 requests a service by sending a message to a subsystem 68, as  
21 represented by the solid arrow between the depicted Win32 client 72 and the  
22 Win32 subsystem 68. The message passes through system services 74 and the  
23 executive 60, which delivers the message to the subsystem. After the subsystem  
24 68 performs the operation, the results are passed to the client 72 in another  
25

1 message, as represented by the arrow between the Win32 subsystem and the  
2 Win32 client 72.

3 In Windows NT, shareable resources, such as files, memory, processes and  
4 threads, are implemented as "objects" and are accessed by using "object services."  
5 As is well known in the art, an "object" is a data structure whose physical format is  
6 hidden behind a type definition. Data structures, also referred to as records or  
7 formats, are organization schemes applied to data so that it can be interpreted and  
8 so that specific operations can be performed on that data. Such data structures  
9 impose a physical organization on the collection of data stored within computer  
10 memory and represent specific electrical or magnetic elements.

11 An "object type," also called an "object class," comprises a data-type,  
12 services that operate on instances of the data type, and a set of object attributes.  
13 An "object attribute" is a field of data in an object that partially defines that  
14 object's state. An "object service" implements and manipulates objects, usually by  
15 reading or changing the object attributes. The object manager 60a is responsible  
16 for creating, deleting, protecting, and tracking user application visible objects.

17 The Windows NT operating system allows users to execute more than one  
18 program at a time by organizing the many tasks that it must perform into  
19 "processes". The operating system allocates a portion of the computer's resources  
20 to each process and ensures that each process's program is dispatched for  
21 execution at the appropriate time and in the appropriate order. This function is  
22 implemented by the process manager 60c.

23 Processes can be implemented as objects. A process object comprises the  
24 following elements: an executable program; a private address space; system  
25

resources (e.g., communication ports and files) that the operating system allocates to the process as the program executes; and at least one "thread of execution."

A "thread" is the entity within a process that the kernel schedules for execution. As is well known in the art, each thread has an associated "context" which is the volatile data associated with the execution of the thread. A thread's context includes the contents of system registers and the virtual address belonging to the thread's process. Thus, the actual data comprising a thread's context varies as it executes. Periodically, a thread may stop executing while, for example, a slow I/O device completes a data transfer or while another thread is using a resource it needs. Because it would be inefficient to have the processor remain idle while the thread is waiting, a multi-tasking operating system will switch the processor's execution from one thread to another in order to take advantage of processor cycles that otherwise would be wasted. This procedure is referred to as "context switching." When the I/O device completes its data transfer or when a resource needed by the thread becomes available, the operating system will eventually perform another context switch back to the original thread. Because of the speed of the processor, both of the threads appear to the user to execute at the same time. Threads that execute programs in the user mode, such as the server 66, the subsystem 68, the application program 70, and the client program 72, are referred to as user mode threads. Threads that execute only operating system components in kernel mode are referred to as kernel mode threads.

### **The Encryption/Decryption Key(s)**

In the described embodiment, use is made of an encryption/decryption key(s). The illustrated and described key is a randomly-generated key that is used



1 to encrypt information that might be paged out to the paging file. The key is  
2 maintained in main memory and is page-locked so that it cannot be paged out to  
3 the secondary memory. Any generation techniques or generators can be used to  
4 generate the random key. An exemplary technique comprises RSA RC4 which is  
5 used for encryption and decryption using a random 2048 bit (256 byte)  
6 cryptographic key. In the described embodiment, this key is used to both encrypt  
7 and decrypt the pageable information. It is to be understood, however, that one  
8 key might be used to encrypt the information while another key might be used to  
9 decrypt the information. In addition, different keys might be used for different  
10 purposes. For example, each separate process might have its own unique key that  
11 is associated with the process. In addition, identifying information might be  
12 encrypted along with the encrypted information to uniquely identify the process  
13 that is associated with the particular encrypted information. Then, only the  
14 process that is associated with the particular encrypted information can decrypt it.

15 In the described embodiment, the information-protection process takes  
16 place in three separate but related phases—an initialization phase, a protection  
17 phase, and an access phase.

### 18 19 **Initialization Phase**

20 The initialization phase is described with reference to Figs. 3 and 4. Fig. 3  
21 shows an exemplary process address space 100, main memory 102, and second  
22 memory (paging file) 104. Fig. 4 shows a flow diagram that describes steps in an  
23 initialization method in accordance with the described embodiment. It is to be  
24 understood that the described initialization phase constitutes but one way of  
25

1 initializing the protection process and is not intended to limit the scope of  
2 protection afforded by this patent.

3 Step 200 (Fig. 4) begins the boot up process for the computer system.  
4 During boot up, the operating system kernel 62 (Fig. 2) allocates a single page of  
5 non-pageable memory (step 202). In the described embodiment, access to the  
6 allocated page is restricted to the software component that does the  
7 encryption/decryption—here, the kernel 62. Thus, user-mode applications cannot  
8 access this memory page. Step 204 generates an encryption key. An exemplary  
9 encryption key was mentioned above. After the encryption key has been  
10 generated, step 206 stores the encryption key in a non-pageable page (i.e. page-  
11 locked) of main memory 102. Fig. 3 shows an exemplary page-locked key at  
12 102a.

13 The implication of having the key stored in a page-locked page in main  
14 memory is that the key can never be paged out to the paging file. This aspect,  
15 combined with the nature of the key (i.e. a very large random key) provides a  
16 degree of protection that previously was not afforded. For example, if the key  
17 cannot be paged out to the paging file, then it is not susceptible to capture.  
18 Additionally, because access to the key is restricted to only the kernel 62, it is  
19 further insulated from attack. Moreover, even if the memory protected by the key  
20 were to be paged out to the paging file, the sheer size of the key alone would make  
21 breaking the protected memory mathematically infeasible in a reasonable amount  
22 of time. And, because the key is randomly generated each time the computer is  
23 booted up, there is a different key for each new computer session. In addition, in  
24 the event of a power loss, the key will be lost from volatile main memory so that  
25 any information that is encrypted and present in the secondary memory will

effectively be lost as well. All of these factors combine to provide a level of protection that greatly improves on those methods used in the past.

### **Protection Phase**

Once the encryption key has been created and stored as described above, information can now be protected. Fig. 3 shows an example of how this can take place. There, a quantity of information 100a is designated as "sensitive information". Such sensitive information is encrypted with the page-locked key 102a to provide encrypted information 100b. The encrypted information 100b is then placed into main memory 102 in a pageable page. Thus, if and when the pageable page that holds the encrypted information 100b is paged out, it is protected in the paging file 104.

The information-protection process can be initiated at the application level or at the memory manager level.

### **Application-level Initiation**

Fig. 5 shows one way that the information protection process can be initiated at the application level. There, the kernel 62 provides interfaces that are callable by the application to protect and unprotect portions of the main memory. Specifically, the kernel 62 provides a protect memory interface 62a and an unprotect memory interface 62b.

Fig. 6 shows a flow diagram that described steps in a method for protecting information using the architecture illustrated in Fig. 5. When an application has information that it wants to protect, it calls the protect memory interface 62a of kernel 62 (step 208) and specifies the address and the size of the information that

1 is to be protected. This call causes the kernel 62 to access the encryption key (step  
2 210) in the page-locked page of main memory 102 (Fig. 3). Once the kernel  
3 accesses the encryption key, it encrypts the information at the specified address  
4 (step 212). The kernel 62 then returns control to the application.

### 6 **Memory Manager-level Initiation**

7 When the information protection process is initiated at the memory  
8 manager 60e level, the memory manager, as Fig. 5 implies, makes the calls to the  
9 protect memory interface 62a of kernel 62. This approach achieves tighter  
10 integration with the memory manager. This method alleviates the need for  
11 application intervention when dealing with sensitive data.

12 Fig. 7 shows a flow diagram that describes steps in a protection process that  
13 is initiated at the memory manager level. In the described embodiment, each page  
14 in the main memory 102 (Fig. 3) that contains information that is to be protected is  
15 designated with a designator (step 214) that can be recognized by the memory  
16 manager. In the described embodiment, the designation is an attribute that can be  
17 set on the page when it is allocated. For example, an attribute "page\_secure" can  
18 be associated with each page and, when set, tells the memory manager 60e that it  
19 must first be encrypted before being paged out. Any time the memory manager  
20 60e handles a page with this designation (i.e. when it is in the process of paging  
21 the page out to the paging file), it recognizes the designation (step 216) and,  
22 responsive thereto, calls the kernel's protect memory interface 62a (step 218).  
23 This call causes the kernel to encrypt the information as described above.

### 24 **Access Phase**

The access phase for accessing a page that has been encrypted and paged out to the paging file is described with reference to Figs. 8 and 9. When a process requires use of information that has been encrypted and paged out to the paging file, step 220 (Fig. 9) accesses the encrypted information 100b. Step 222 accesses the page-locked key 102a and uses it to decrypt (step 224) the encrypted information. The decrypted information is then placed in main memory 102 where it can be used by the process.

In the described embodiment, decryption can be initiated at the application level or at the memory manager level. At the application level, the application can call the unprotect memory interface 62b (Fig. 5) of kernel 62. This call causes the kernel to retrieve the information from the paging file, decrypt it using the page-locked key, and then turn the information over to the application or process. At the memory manager level, when the memory manager retrieves an encrypted page from the paging file, it recognizes the designation (i.e. attribute) on the page and calls the unprotect memory interface 62b of the kernel 62. This call causes the kernel to decrypt the information as mentioned above.

When the information is decrypted and placed into main memory, the page into which it is placed should be zeroed as soon as possible to avoid page file exposure. As an added measure of safety, the decrypted information can be placed into a page-locked page to ensure that it does not inadvertently get paged out to the paging file; when the application is finished with the data, it would zero the memory, and then unlock the memory page, assuming it were locked.

1        **Conclusion**

2        The inventive methods and systems provide degrees of protection for  
3 information that might be paged out to a paging file which were heretofore  
4 unavailable. The size, nature and handling of the key(s) that is (are) used for  
5 encryption greatly increases the protection of pageable information.

6        Although the invention has been described in language specific to structural  
7 features and/or methodological steps, it is to be understood that the invention  
8 defined in the appended claims is not necessarily limited to the specific features or  
9 steps described. Rather, the specific features and steps are disclosed as preferred  
10 forms of implementing the claimed invention.